



# UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE  
United States Patent and Trademark Office  
Address: COMMISSIONER FOR PATENTS  
P.O. Box 1450  
Alexandria, Virginia 22313-1450  
www.uspto.gov

*m*

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
09/884,467	06/20/2001	Jeffrey A. Bedell	53470.003030	9623

21967 7590 02/20/2008  
HUNTON & WILLIAMS LLP  
INTELLECTUAL PROPERTY DEPARTMENT  
1900 K STREET, N.W.  
SUITE 1200  
WASHINGTON, DC 20006-1109

EXAMINER

WASSUM, LUKE S

ART UNIT PAPER NUMBER

2167

MAIL DATE DELIVERY MODE

02/20/2008

PAPER

**Please find below and/or attached an Office communication concerning this application or proceeding.**

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

---

Commissioner for Patents  
United States Patent and Trademark Office  
P.O. Box 1450  
Alexandria, VA 22313-1450  
[www.uspto.gov](http://www.uspto.gov)

BEFORE THE BOARD OF PATENT APPEALS  
AND INTERFERENCES

**MAILED**

**FEB 20 2008**

**Technology Center 2100**

Application Number: 09/884,467  
Filing Date: June 20, 2001  
Appellant(s): BEDELL ET AL.

---

Brian M. Buroker  
For Appellant

**EXAMINER'S ANSWER**

This is in response to the appeal brief filed 7 December 2007 appealing from the Office action mailed 2 August 2007.

**(1) Real Party in Interest**

A statement identifying by name the real party in interest is contained in the brief.

**(2) Related Appeals and Interferences**

The examiner is not aware of any related appeals, interferences, or judicial proceedings which will directly affect or be directly affected by or have a bearing on the Board's decision in the pending appeal.

**(3) Status of Claims**

The statement of the status of claims contained in the brief is correct.

**(4) Status of Amendments After Final**

No amendment after final has been filed.

**(5) Summary of Claimed Subject Matter**

The summary of claimed subject matter contained in the brief is correct.

**(6) Grounds of Rejection to be Reviewed on Appeal**

The appellant's statement of the grounds of rejection to be reviewed on appeal is correct.

**(7) Claims Appendix**

The copy of the appealed claims contained in the Appendix to the brief is correct.

**(8) Evidence Relied Upon**

**Hahn et al.** "Capacity Planning for Business Intelligence Applications: Approaches and Methodologies", document SG24-5689-00, November 2000.

**IBM[1]** "OS/390 Resource Measurement Facility Report Analysis", document SC28-1950-05, 6th Edition, September 2000, pp. 1-1 through 1-7, 2-240 through 2-262 and 5-1 through 5-52.

**(9) Grounds of Rejection**

The following ground(s) of rejection are applicable to the appealed claims:

Claim 28 is rejected under 35 U.S.C. 112, first paragraph, as failing to comply with the written description requirement. The claim(s) contains subject matter which was not described in the specification in such a way as to reasonably convey to one skilled in the relevant art that the inventor(s), at the time the application was filed, had possession of the claimed invention.

Claims 1-8, 10-17, 19-26 and 29 are rejected under 35 U.S.C. 102(a) as being anticipated by **IBM** ("The Business Intelligence Infrastructure on S/390 Accessing DB2 on OS/390") as evidenced by **Hahn et al.** ("Capacity Planning for Business Intelligence Applications: Approaches and Methodologies") and **IBM[1]** ("OS/390 Resource Measurement Facility Report Analysis").

Claims 9, 18 and 27 are rejected under 35 U.S.C. 103(a) as being unpatentable over **IBM** ("The Business Intelligence Infrastructure on S/390 Accessing DB2 on OS/390") as evidenced by **Hahn et al.** ("Capacity Planning for Business Intelligence Applications: Approaches and Methodologies") and **IBM[1]** ("OS/390 Resource

Measurement Facility Report Analysis") as applied to claims 1-8, 10-17, 19-26 and 29 above, and further in view of **Hayes** (U.S. Patent Application Publication 2002/0046204).

These rejections are set forth in a previous Office action, mailed 5 March 2007.

For the convenience of the Honorable Board of Appeals, the rejection of representative claim 1 is reproduced herein.

Regarding claim 1, **IBM** teaches a computer-implemented method for capturing at least one statistic or data regarding performance operation of a business intelligence reporting system that generates business intelligence reports based on requests submitted to perform analysis of data contained in a database as claimed, the process comprising the steps of:

- a) gathering at least one statistic or data related to the performance operation of the reporting system while the reporting system is operating (see disclosure of the System Management Facility which measures various aspects of the work running in the system, such as measures of system resources utilized;

for instance, CPU time, storage, I/O devices, and service units, **Hahn et al.**

section 5.1.1 System Management Facility (SMF), pages 65-66 et seq.);

b) analyzing the at least one statistic or data (see disclosure of the configuration

of definitions and criteria for workflow exceptions in the Resource

Measurement Facility, **IBM[1]** pages 2-251 through 2-256, whereby

exception conditions are defined); and

c) generating at least one output based on the gathered at least one statistic or

data (see disclosure of the Resource Management Facility, which can

generate reports summarizing key resources, including the CPU Activity

Report and the Workload Activity Report, **Hahn et al.** section 5.1.2

Resource Management Facility (RMF), page 66 et seq.), wherein the at least

one output includes an alert if the analysis of the at least one statistic or

data indicates that a condition has occurred (see disclosure of the

configuration of alerts in the Resource Measurement Facility, **IBM[1]** page

2-255, top; see also drawing Figures 2-114 and 2-115, on pages 2-252 and 2-

256 respectively, showing the configuration of alerts).

#### **(10) Response to Argument**

This Examiner's Answer will address the Appellant's arguments in the order in which they appear in the appeal brief.

**A. Issue 1**

**The Rejection of Claim 28 under 35 U.S.C. § 112, First Paragraph**

Regarding the Appellants' arguments on pages 6-8 of the Appeal Brief, the examiner finds these arguments persuasive. The rejection of claim 28 under 35 U.S.C. § 112, first paragraph is therefore withdrawn.

**B. Issue 2**

**The Rejection of Claims 1-8, 10-17, 19-26 and 29 under 35 U.S.C. § 102(a)**

Regarding claims 1, 10 and 19, the Appellants argue that (1) the prior art of record fails to disclose the gathering of statistics or data related to the performance/operation of a reporting system, that (2) it is quite possible to have poor performance on a report while the underlying operating system, such as OS/390, is



performing well and has low utilization, and that (3) the system management facility reports disclosed by the prior art of record does not teach "capturing at least one statistic or data regarding performance operation of a business intelligence reporting system that generates business intelligence reports based on requests submitted to perform analysis of data.

In response, the examiner presents the following arguments.

Regarding argument (1) [that the prior art of record fails to disclose the gathering of statistics or data related to the performance/operation of a reporting system], the examiner respectfully disagrees.

The essential disagreement on this issue seems to be whether the statistics disclosed in the prior art of record qualify as the claimed 'statistics/data related to the operation/performance of a report system'.

The Appellants' specification, at page 1, lines 12-15, disclose that the reporting system is a business intelligence system:

**Background of the Invention**

Decision support systems have been developed to efficiently retrieve selected information from data warehouses. One type of decision support system is known as an on-line analytical processing system. In general, business intelligence, OLAP or reporting systems analyze data from various perspectives and support complex analyses against large input data sets.

The prior art of record discloses such a Business Intelligence system. For instance, the **Hahn et al.** reference discloses capacity planning for Business Intelligence applications, as is reflected in the title of the reference.

In section 1.2, beginning on page 6, the reference includes some details of the disclosed Business Intelligence infrastructure, including three different architectures in section 1.2.1.

In particular, on page 7 in section 1.2.1, **Hahn et al.** discloses the components of the Business Intelligence system:

### **1.2.1 The Business Intelligence system architectures**

The Business Intelligence system consists of all the software that performs the processing of data described in the previous section, and the physical hardware and operating system upon which the software runs. When deploying BI applications, as with other applications, there are three logical processing layers, as shown in Figure 4 on page 8:

- *The data layer* - This layer supports all processes that move, store, format, and change the raw data. This includes all of the processes of data collection and maintenance of the data warehouse and data marts, as well as the processing of Structured Query Language (SQL).
- *The application logic layer* - This layer supports secured access to the BI application and data, the creation of SQL, and post processing of retrieved data before presentation to the business user. This may include security authorization and validation, generation of SQL, and manipulation to complete business rules processing and analysis.
- *The presentation logic layer* - This layer directly interacts with the business user. It includes the graphical user interface (GUI) to develop business queries and to present results in multiple formats, i.e. texts, charts, and graphics.

Most notably, it is disclosed that the Business Intelligence system consists of all the software that performs the processing of data, as well as the physical hardware and operating system upon which the software runs.

This expansive view of the components of a business intelligence reporting system is supported by the Appellants' own disclosure as well. For instance, regarding the statistics to be captured, it is disclosed in the specification on pages 3-4 that

Relevant statistics and data may include, but are not limited to, overall system usage, individual user activity, usage data about reports, usage data about one or more objects used to operate and interact with the reporting system, database usage data (*e.g.*, workload, response times, the number of connections to the database, *etc.*), and concurrency information (*e.g.*, the number of concurrent users, the number of concurrent requests, the number of requests processed concurrently, the number of requests in the queue at any given time, *etc.*). Relevant statistics and data may also include server statistics (*e.g.*, amount of memory, computer processing unit (CPU) and input/output resources utilized at a given point) and metadata usage (*e.g.*, the number of requests for metadata, the types of requests for metadata, *etc.*). Other relevant statistics and data may also be used.

The Appellants disclose that 'relevant statistics' include, amongst other things, 'overall system usage, usage data about one or more objects used to operate and interact with the reporting system, database usage data (*e.g.*, workload, response times, the number of connections to the database, *etc.*)', as well as 'server statistics (*e.g.*, amount of memory, computer processing unit (CPU) and input/output resources utilized at a given point)'. This disclosure indicates that the type of statistics that are 'relevant to the performance of a business intelligence reporting system' go beyond those argued by the Appellants in their Brief, and similarly that a business intelligence reporting system encompasses more than the narrow definition argued by the Appellants.

Regarding argument (2) [that it is quite possible to have poor performance on a report while the underlying operating system, such as OS/390, is performing well and has low utilization], the examiner respectfully responds that this point, while having merit, is not relevant to the claim limitations at issue.

The claim language calls for statistics or data to be gathered. The statistics or data are related to the performance of a business intelligence reporting system. As argued above, both the prior art of record and the Appellants' own disclosure indicate that the claimed business intelligence reporting system encompasses more than is argued by the Appellants in their Brief. The specification discloses that relevant statistics would include server statistics and database usage data. This means that data regarding the database and operating system, such as is disclosed in prior art of record, would constitute the claimed statistics or data.

Even were the Appellants' arguments regarding the narrow view of what constitutes the claimed business intelligence reporting system to be adopted, the claim

language is broad enough to be anticipated the statistics and data disclosed by the prior art of record.

The claim calls for statistics or data to be 'related to the performance of the reporting system'. Certainly, the performance of the operating system and the performance of the database management system will have some relationship to the performance of the overall business intelligence reporting system. Were there to be a serious performance problem with either the operating system or the database management system, this would surely be reflected in the performance of the business intelligence reporting system as a whole. As such, statistics or data such as is disclosed in the prior art of record is surely 'related' to the performance of the overall business intelligence reporting system, and anticipates the claims.

Regarding argument (3) [that the system management facility reports disclosed by the prior art of record does not teach "capturing at least one statistic or data regarding performance operation of a business intelligence reporting system that generates business intelligence reports based on requests submitted to perform analysis of data], the examiner respectfully disagrees.

The Appellants are apparently arguing that the limitation at issue should be interpreted as requiring a user to request 'analysis of data' before any statistics and/or data regarding performance operation of a business intelligence reporting system is captured.

The examiner believes that the phrase 'based on requests submitted to perform analysis of data in a database' instead refers to the user requests to generate business intelligence reports.

The examiner points out that this interpretation must be assumed correct, since if the 'user requests to perform analysis of data in a database' was referring to a user request to perform analysis of system performance statistics, this operation would be impossible, since there would be no data to be analyzed, that is, statistics or data regarding performance operation of a business intelligence reporting system.

Put another way, the Appellants argue that no statistics are collected until the user submits a request to analyze the statistics. Clearly, this cannot be the correct interpretation of the limitation at issue. The correct interpretation is that the statistics gathered are based upon user requests to perform analysis of the business intelligence data in the system.

Further support for this interpretation is provided by the Appellants' own specification, for instance, on page 11:

The present invention provides for capturing relevant statistics and data to enable profiling and/or understanding of current and historical activity of the reporting system.

Statistics may be stored, where storage methods disk based, file system based, database based, object oriented database based (*e.g.*, relational databases) or others, such as, but not limited to, NT Performance Monitor™ and NT Events Log™. Stored statistics may enable additional analysis of a reporting system activity, including identifying trends. Analyzing historical report system activity may include the ability to utilize decision support techniques and/or software to facilitate analysis.

It seems clear that in order for 'historical activity' to be analyzed, and to allow for 'identifying trends', the statistics regarding performance operation of a business intelligence reporting system must be collected on a regular basis, and not merely in response to a user request to analyze the statistical data that would not even exist were the Appellant's interpretation of the claim limitation to be adopted.



For these reasons, the examiner maintains that the rejections of claims 1-8, 10-17, 19-26 and 29 are proper, and should be sustained.

### **C. Issue 3**

#### **The Rejection of Claims 9, 18 and 27 under 35 U.S.C. § 103(a)**

Regarding claims 9, 18 and 27, the Appellants argue that (1) the prior art of record fails to teach tuning a report system, but instead are directed to operating systems, that (2) the **Hayes** reference does not disclose "performing automated tuning of the reporting system based on the at least one output", and that (3) the **Hayes** reference fails to teach the automated tuning of the system based upon an output which is "based upon requests submitted to perform analysis of data in the database".

In response, the examiner presents the following arguments.

Regarding argument (1) [that the prior art of record fails to teach tuning a report system, but instead are directed to operating systems], the examiner respectfully

disagrees. This argument has been addressed above with regard to the rejections under 35 U.S.C. § 102(a).

Regarding argument (2) [that the **Hayes** reference does not disclose "performing automated tuning of the reporting system based on the at least one output"], the examiner points out that the rejection of record does not allege that **Hayes** teaches this limitation.

As discussed above, the **Hahn et al.** and **IBM[1]** references are relied upon for their teaching of the capturing of statistics related to the performance operation of a business intelligence reporting system. These references also teach, as recited in the rejections of record, that these statistics are used in tuning the settings of the business intelligence system in order to optimize performance, specifically in the **Hahn et al.** reference in section 5.1.2 on page 66.

### **5.1.2 Resource Management Facility (RMF)**

The Resource Monitoring Facility (RMF) is also a component of the OS/390 operating system. RMF contains three separate monitors that can be operated to collect and write data into SMF records (types 70-79). These records are created and stored in the SMF datasets, and eventually are offloaded. The RMF postprocessor routine can then process these RMF records to create comprehensive reports of either a single-system or sysplex scope.

The RMF reports generated by the postprocessor are used for either capacity planning or performance tuning of an OS/390 system. For the capacity planning activity, only three reports are necessary:

- The Summary Report summarizes the key resources for an OS/390 image.
- The CPU Activity Report provides an overview of activity for all processors belonging to this OS/390 image.
- The Workload Activity Report shows workload-related data on resource consumption with different levels of detail. It provides a summary by performance group (service class in goal mode), by reporting groups, and for the entire system.

To create these postprocessor reports, code the batch job using SMF datasets as input. Refer to *RMF Users Guide* for more information. Later in this chapter we discuss each worksheet, as well as the exact reports and fields for input for these reports. For additional information on RMF reports, see either *RMF Performance Management Guide* or *RMF Report Analysis*.

Also relevant are the disclosures in section 5.1.3 through 5.1.5.

### 5.1.3 DB2 Performance Monitor

IBM DATABASE 2 Performance Monitor (DB2 PM) is a tool for analyzing and tuning the performance of DB2 and DB2 applications. It runs on an OS/390 server and has a real-time online monitor which runs on an OS/2 or a Windows NT workstation. It includes a history facility to view recent events,

and a wide variety of reports for in-depth analysis. DB2 PM V6 has been included as a feature of the DB2 UDB for OS/390 V6 product.

DB2 generates data about its own performance, called *instrumentation data*, but it does not provide any reporting facilities for analyzing this data. The Online Monitor of DB2 PM allows you to view an active DB2 subsystem and identify performance problems online.

For a long-term view of DB2 performance, use the DB2 PM Batch reporting capabilities, which also provide an historical view of DB2 performance. When you are performing capacity planning, DB2 PM can be used to measure an application's performance, its resource usage, and the effect an application can have on the system.

### 5.1.4 How performance data is generated

The DB2 trace facility, also called the DB2 instrumentation facility, gathers information about data and events in the database. After DB2 has collected and externalized this data, DB2 PM reads it and generates reports, graphs, and data sets from it.

When starting the DB2 trace facility, specify the performance information you want to collect. The data is recorded in different record types, called instrumentation facility component identifiers (IFCIDs). These records are grouped into *classes*, and the classes are grouped into *types*, which are the broadest categorization of performance data. The types used by DB2 PM Batch are:

- Statistics
- Accounting
- Audit
- Performance
- Global

SMF is the default destination of trace output for Statistics and Accounting information. All DB2 trace data is found in SMF record types 100 through 102.

Traces have a performance overhead, so carefully consider which traces to run. IBM recommends regularly collecting accounting class 1 and 3 and statistics class 1, 3, and 4 data. You should also consider collecting accounting class 2 data, which provides important information about DB2 times.

### 5.1.5 How DB2 PM reports the data

DB2 PM generates reports, data sets, graphs, and logs. The output is controlled by means of commands. The output is grouped into *report sets*, each associated with a particular type of data. The report sets most commonly used for capacity planning are *Statistics* and *Accounting*. Statistics summarizes system-wide performance data, whereas Accounting summarizes information for particular applications. Both report sets are used to determine the efficiency of the subsystem or application.

Some forms of DB2 PM Batch output we are interested in are:

- Traces

Traces show individual DB2 events, for example, for a particular thread. All events are listed individually, usually in the order of occurrence.

- Reports

Reports show these events summarized by identifiers, such as primary authorization ID or plan name.

- Data sets

Formatted data can be stored in data sets suitable for loading into DB2 tables. The tables can then be easily queried to produce customized reports using a reporting tool such as IBM Query Management Facility (QMF).

These tables are optional, but we strongly recommend having historical performance data loaded into DB2 tables. DB2 PM provides the statements for creating the tables, loading the tables, and sample queries. The data can be used for capacity planning as well as performance and tuning. Storing the information in DB2 data sets simplifies saving the data for use for future analysis or problem determination. We will simply refer to these tables as the performance database.

For additional information refer to *DB2 Performance Monitor for OS/390 Batch User's Guide* and *DB2 Performance Monitor for OS/390 Report Reference, Volume I* and *Volume II*.

The Hayes reference is relied upon merely to teach that the use of captured performance statistics in tuning system settings to optimize performance can be done as an automated process, instead of being performed manually.

This is taught in paragraph [0084]:

**[0084]** Turning now to the Figures, FIGS. 1-4, depicted therein is a flowchart for an embodiment of the present invention. FIG. 1 depicts the initial steps of a bufferpool tuning algorithm, based on the next higher size pool, if available, within the past thirty days, based on actual usage statistics, and adjusts the bufferpool accordingly.

Therein, the reference teaches an algorithm that automatically tunes a bufferpool based upon actual usage statistics. The rejection of record relies upon the reference for the automated aspect of the tuning of the system based upon usage statistics.

Regarding argument (3) [that the Hayes reference fails to teach the automated tuning of the system based upon an output which is "based upon requests submitted to perform analysis of data in the database"], the examiner respectfully disagrees.

This argument is actually an extension of argument (3) presented with regard to the rejections of record under 35 U.S.C. § 102(a), which have been addressed by the

examiner above. In short, the capturing of statistics cannot be properly interpreted as occurring only in response to a user request to perform analysis of those same statistics, since were this to be the case, there would be no statistics to analyze.

For these reasons, the examiner maintains that the rejections of claims 9, 18 and 27 are proper, and should be sustained.

**(11) Related Proceeding(s) Appendix**

No decision rendered by a court or the Board is identified by the examiner in the Related Appeals and Interferences section of this examiner's answer.

*Conclusion*

Claims 1-8, 10-17, 19-26 and 29 are properly rejected under 35 U.S.C. § 102(a) as being anticipated by **IBM** ("The Business Intelligence Infrastructure on S/390 Accessing DB2 on OS/390") as evidenced by **Hahn et al.** ("Capacity Planning for Business Intelligence Applications: Approaches and Methodologies") and **IBM[1]** ("OS/390 Resource Measurement Facility Report Analysis").

Claims 9, 18 and 27 are properly rejected under 35 U.S.C. 103(a) as being unpatentable over **IBM** ("The Business Intelligence Infrastructure on S/390 Accessing DB2 on OS/390") as evidenced by **Hahn et al.** ("Capacity Planning for Business Intelligence Applications: Approaches and Methodologies") and **IBM[1]** ("OS/390 Resource Measurement Facility Report Analysis") as applied to claims 1-8, 10-17, 19-26 and 29 above, and further in view of **Hayes** (U.S. Patent Application Publication 2002/0046204).



In light of the foregoing arguments, the Examiner respectfully requests the Honorable Board of Appeals to sustain the rejections.


Respectfully submitted,




Luke S. Wassum  
Primary Examiner  
Art Unit 2167

Conferees:  
John R. Cottingham  
Eddie C. Lee

lsw  
12 February 2008



**EDDIE C. LEE**  
**SUPERVISORY PATENT EXAMINER**



**JOHN COTTINGHAM**  
**SUPERVISORY PATENT EXAMINER**  
**TECHNOLOGY CENTER 2100**